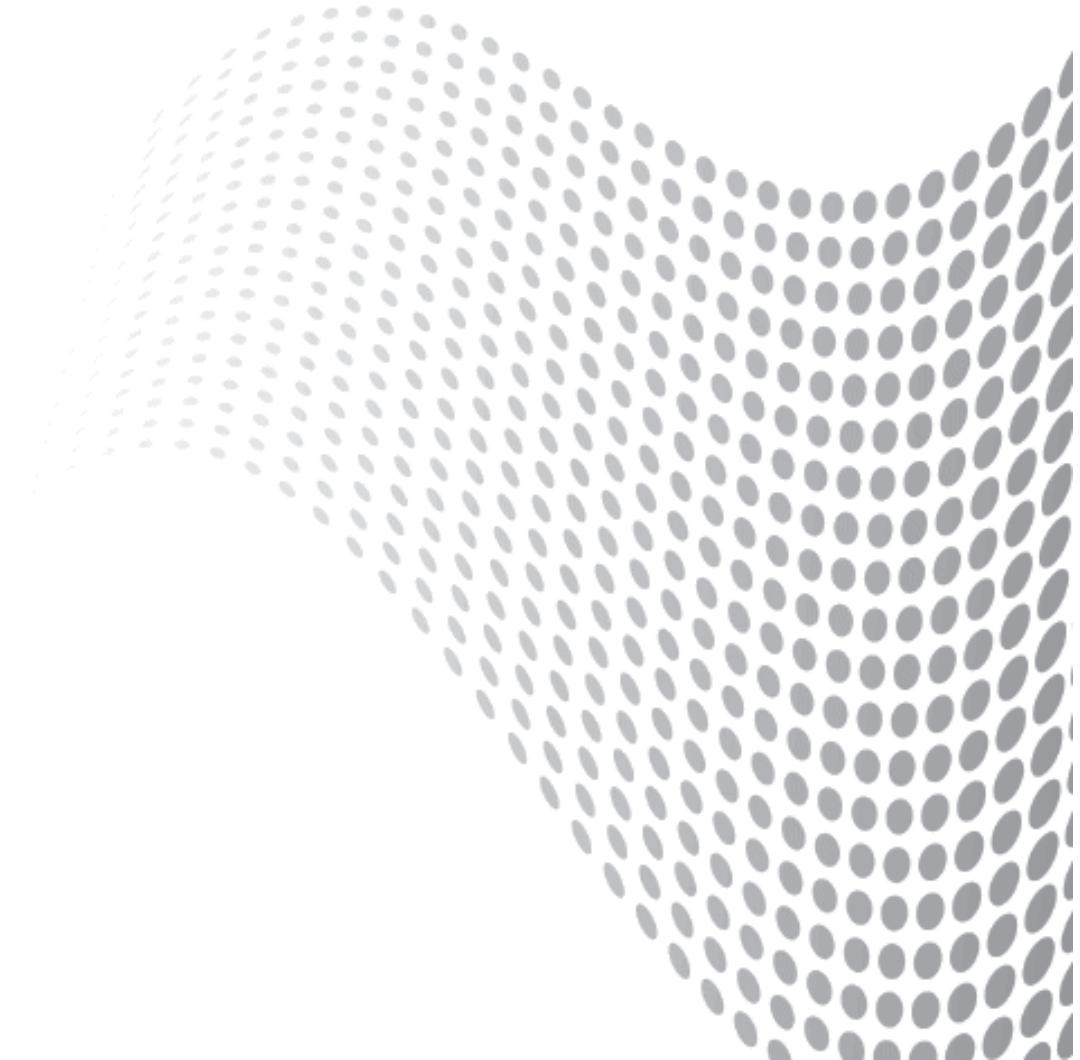




FAST Enterprise Search Platform

version:5.3

Troubleshooting Guide



Copyright

Copyright © 1997-2009 by Fast Search & Transfer ASA (“FAST”). Some portions may be copyrighted by FAST’s licensors. All rights reserved. The documentation is protected by the copyright laws of Norway, the United States, and other countries and international treaties. No copyright notices may be removed from the documentation. No part of this document may be reproduced, modified, copied, stored in a retrieval system, or transmitted in any form or any means, electronic or mechanical, including photocopying and recording, for any purpose other than the purchaser’s use, without the written permission of FAST. Information in this documentation is subject to change without notice. The software described in this document is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Trademarks

FAST ESP, the FAST logos, FAST Personal Search, FAST mSearch, FAST InStream, FAST AdVisor, FAST Marketrac, FAST ProPublish, FAST Sentimeter, FAST Scope Search, FAST Live Analytics, FAST Contextual Insight, FAST Dynamic Merchandising, FAST SDA, FAST MetaWeb, FAST InPerspective, NXT, FAST Unity, FAST Radar, RetrievalWare, AdMomentum, and all other FAST product names contained herein are either registered trademarks or trademarks of Fast Search & Transfer ASA in Norway, the United States and/or other countries. All rights reserved. This documentation is published in the United States and/or other countries.

Sun, Sun Microsystems, the Sun Logo, all SPARC trademarks, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Netscape is a registered trademark of Netscape Communications Corporation in the United States and other countries.

Microsoft, Windows, Visual Basic, and Internet Explorer are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Red Hat is a registered trademark of Red Hat, Inc.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

AIX and IBM Classes for Unicode are registered trademarks or trademarks of International Business Machines Corporation in the United States, other countries, or both.

HP and the names of HP products referenced herein are either registered trademarks or service marks, or trademarks or service marks, of Hewlett-Packard Company in the United States and/or other countries.

Remedy is a registered trademark, and Magic is a trademark, of BMC Software, Inc. in the United States and/or other countries.

XML Parser is a trademark of The Apache Software Foundation.

All other company, product, and service names are the property of their respective holders and may be registered trademarks or trademarks in the United States and/or other countries.

Restricted Rights Legend

The documentation and accompanying software are provided to the U.S. government in a transaction subject to the Federal Acquisition Regulations with Restricted Rights. Use, duplication, or disclosure of the documentation and software by the government is subject to restrictions as set forth in FAR 52.227-19 Commercial Computer Software-Restricted Rights (June 1987).

Contact Us

Web Site

Please visit us at: <http://www.fastsearch.com/>

Contacting FAST

FAST
Cutler Lake Corporate Center
117 Kendrick Street, Suite 100
Needham, MA 02492 USA
Tel: +1 (781) 304-2400 (8:30am - 5:30pm EST)
Fax: +1 (781) 304-2410

Technical Support and Licensing Procedures

Technical support for customers with active FAST Maintenance and Support agreements, e-mail: fasttech@microsoft.com

For obtaining **FAST licenses or software**, contact your FAST Account Manager or e-mail: fastcsrv@microsoft.com

For **evaluations**, contact your FAST Sales Representative or FAST Sales Engineer.

Product Training

E-mail: fastuniv@microsoft.com

To access the FAST University Learning Portal, go to: <http://www.fastuniversity.com/>

Sales

E-mail: sales@fastsearch.com

Contents

Preface.....	ii
Copyright.....	ii
Contact Us.....	iii
Chapter 1: Gathering support information.....	9
Gathering support information.....	10
Reporting issues.....	10
Retrieving product version information.....	10
Retrieving platform/operating systems information.....	10
Environment variables information.....	11
License status.....	11
Machine load and memory.....	11
Disk space and cpu usage.....	11
IP and network addresses.....	12
Ports allocation.....	12
Process status.....	12
Documents status.....	12
Chapter 2: infotool.....	15
Running infotool.....	16
infotool options.....	16
infotoolConfig.xml.....	16
Collectors.....	17
InfotoolComplete.....	17
ChecksumGenerator.....	17
ConfigurationReport.....	18
CoreDetector.....	18
CrawlerConfig.....	18
FASTVersion.....	18
ConfigZipper.....	18
FreeSpace.....	19
Logs.....	19
OperatingSystem.....	20
SystemMapper.....	20
HardwareProfile.....	20
DocumentCounter.....	20
Eventlog.....	20
LocaleAndTime.....	21

ViewCollector.....	21
--------------------	----

Chapter 3: Log file information.....23

About log files.....	24
Log file types.....	24
Guidelines for using log files.....	24
Log file information for individual modules.....	24
Configuration server.....	25
Content API or SDK.....	25
Crawler.....	25
Document processor(s).....	26
esp4jtool-dictman dictionary management tool.....	27
File traverser.....	27
License manager.....	27
Node controller.....	28
QRServer.....	28
Real-time search engine/Real-time search indexer.....	30
Storage service.....	30
WebAnalyzer.....	30

Chapter 4: Known issues and troubleshooting tips.....33

Document processor issues.....	34
Optimize document processing performance.....	34
Unusual behavior during document processing.....	34
Documents skipping stages.....	34
PDF documents missing attributes in result set.....	35
Lemmatization tips and issues.....	35
Lemmatization tips.....	35
Add a language for lemmatization.....	35
Change lemmatization for a language.....	36
Set up default query lemmatization.....	36
Synonyms issues.....	37
Index-side synonyms do not work.....	37
Index-side synonyms do not work for all documents.....	37
Index-side synonyms do not work for all queries.....	37
Query-side synonyms seen in the system log.....	37
Real-time search engine tips and issues.....	38
Real-time search engine tips.....	38
Systems unable to index.....	38
Identifying and recovering from search node failure.....	39
QRServer issues.....	39
QRServer and fdispatch not connected.....	39
QRServer timed out.....	40
License manager tips and issues.....	40

License manager tips.....	40
Components not initializing.....	40
QRServer processing is limited.....	40
Disk configuration tips.....	41
Network tips.....	41
Query transmission is slow.....	42
Machine configuration issues	42
SBC issues.....	42
SBC query reports cannot be generated.....	42
Chapter 5: Log levels and log messages.....	43
Logserver configuration.....	44
Determining the destination host for log messages.....	44
Log levels and log messages for individual modules.....	45
Admin server log levels and messages.....	45
Configuration server log levels and messages.....	45
Content distributor log levels and messages.....	48
Crawler error codes and messages.....	52
File traverser log levels and messages.....	52
License error messages.....	55
Postprocess error codes and messages.....	55
Node controller log levels and messages.....	55
Chapter 6: Processes overview.....	57
Processes descriptions.....	58

Chapter

1

Gathering support information

Topics:

- *Gathering support information*
- *Reporting issues*
- *Retrieving product version information*
- *Retrieving platform/operating systems information*
- *Environment variables information*
- *License status*
- *Machine load and memory*
- *Disk space and cpu usage*
- *IP and network addresses*
- *Ports allocation*
- *Process status*
- *Documents status*

This chapter describes the type of operational information that is critical in providing a fast and complete response from FAST, and how to gather this information.

Gathering support information

Before you contact FAST Technical Support with an issue, gathering specific information will help to identify and solve the issue in a timely manner. This section describes which type of information you need to submit to FAST Technical Support and how you gather this information.



Note: infotool can be used to gather much of this information.

Reporting issues

When reporting operational issues, the following list of information is critical in providing a fast and complete response:

- Product Version - the version of FAST ESP you are running and any patches that have been applied.
- Platform - the platform/operating system on which you are running.
- Overall System - the type of system, production or development, that is encountering the issue to help determine the urgency of the issue.
- Log Files - any necessary log files for background information.
- Issue Description - a detailed description of the issue.
- Issue Reproduction - detailed steps to reproduce the issue.

Retrieving product version information

Use either of the following actions to locate the version of product that you are running:

- Click on **System Overview** in the FAST ESP Administrator Interface and note the **Build-number**.
- Examine the \$FASTSEARCH/etc/VERSION.xml file.

Retrieving platform/operating systems information

FAST ESP versions maintain the same directory structure, configuration files and binaries across all platforms. The binaries are compiled for each platform.



Note: Some paths may be described as UNIX paths. Windows users can modify them by changing the environmental variable from \$FASTSEARCH to %FASTSEARCH% and reversing the slashes in the path. For example: \$FASTSEARCH/bin on UNIX is equivalent to %FASTSEARCH%\bin on Windows.

To obtain the existing operating system information:

- For Windows, verify whether you are running on Microsoft Windows 2000 or 2003.
- For UNIX, log on as a UNIX user and type `uname -a` to obtain operating system, hardware architecture and operating system version information.

For Linux systems, log on as a UNIX user and type `uname -a` to obtain operating system and hardware architecture information. Report the content in the \$FASTSEARCH /etc/issue files to include operating system version information.

Environment variables information

This section describes system and machine environment variables.

System Environment

- The `FASTSEARCH` variable must point to the root of the FAST ESP installation directory.
- `LM_LICENSE_FILE` variable must point to the installed license file, `$FASTSEARCH/etc/fastsearch.lic`.
- `LD_LIBRARY_PATH` (Solaris/Linux) environment variable provides the location of the shared libraries and must point to `$FASTSEARCH/lib`.

Machine Environment

- For Windows, right-click on **My Computer** and select **Properties**. Information for the system is listed under the **General** tab.
- For UNIX (Linux, Solaris), type in the following commands: `uname -a` and `ulimit -a`

License status

Your FAST ESP product may include both standard and optional features. Optional features are enabled with individual license keys. If you have not purchased these optional features, they will not be enabled in your FAST ESP installation.

 **Note:** If your license server fails, FAST ESP will run an additional four days. However, you will not be able to restart FAST ESP.

If you encounter issues with the license file:

- Check that the host name of the License Manager is correct in the SERVER line of the license file. The license file can be found in the `$FASTSEARCH/etc` directory of the server running the License Manager (usually the administration node).
- Check to make sure that the license has not expired (if applicable). This should be described in the header text of the license file.

If these suggestions do not solve the issue, then go to the following files and report your findings:

```
$FASTSEARCH/etc/fastsearch.lic
$FASTSEARCH/var/log/lmgrd/lmgrd.scrap    (log file for License Manager)
$FASTSEARCH/var/log/*.log files
```

Machine load and memory

For Windows, run `msinfo32` and provide the **System Summary** information.

For UNIX (Linux), use the `free` and `vmstat` commands to get memory statistics.

For UNIX (Solaris), use the `vmstat` command to get memory statistics.

Disk space and cpu usage

For Windows, go into **Task Manager** and select the **Processes** tab to view CPU and memory usage.

For UNIX, use the `iostat` and `top` commands.

IP and network addresses

For Windows, use the `ipconfig` command.

For UNIX:

- Check the `hostname` and `/etc/hosts` and `/etc/resolv.conf` files.
- Use the `ifconfig -a` command.
- Use the `nslookup` command for DNS check.

Ports allocation

Select **System Overview** in the FAST ESP Administrator Interface to view the main ports used by the FAST ESP modules.

For additional ports allocation information, use one of the following commands:

- For Windows, use the `netstat` command.
- For UNIX, use the `netstat -p` command.

Process status

Select **System Overview** in the FAST ESP Administrator Interface to view process status.

For additional process status information, use one of the following commands:

- For Windows (OS level): Go into **Task Manager** and select the **Processes** tab to view the status of all processes.
- For UNIX (OS level): Use the `ps` command.
- For FAST ESP: Use the `$FASTSEARCH/bin/nctrl sysstatus` command. This command applies to both Windows and UNIX platforms.

Documents status

If documents are not being indexed but have been crawled, it is important to determine where in the process the document gets dropped. It is possible to trace a document to see how it has been processed.

To trace a document's progress:

1. Execute the `doctrace` tool from the command prompt.

```
$FASTSEARCH/bin/psctrl doctrace on
```

This tool enables document tracing as documents flow through the pipeline.

2. Add a single document that contains the data in question to a collection.
3. Execute the following to monitor how the document is being modified by the pipeline:

```
$FASTSEARCH/bin/doclog -a
```

4. Check the log files for any possible errors. In such cases, there will be a mismatch between the number of documents in the content distributor (accessed from the **Collection Overview** selection) and the number of documents in the index.

Chapter

2

infotool

Topics:

- [*Running infotool*](#)
- [*infotool options*](#)
- [*infotoolConfig.xml*](#)
- [*Collectors*](#)

As described in the *Gathering support information* chapter, gathering information before you contact FAST Technical Support will help to identify and solve the issue in a timely manner. This chapter describes infotool, a tool to help you collect information about a FAST ESP installation, in order to help FAST Technical Support diagnose problems with the installation.

Running infotool

Run infotool without options if you are collecting information to submit to FAST Technical Support.

To run infotool, execute the default collection pipeline.

1. Run the following:

For UNIX:

```
$FASTSEARCH/bin/infotool
```

For Windows:

```
%FASTSEARCH%\bin\infotool.exe
```

infotool options

```
$FASTSEARCH/bin/infotool <options...>
```

Option	Description
-h	Display usage.
-v	Print version.
-t	Test a pipeline configuration.
-p	Select which pipeline to run from InfotoolConfig.xml. Example: infotool -p configreport
-c	Clock the running time of each collector.
-d	Enable debug logging.
-r	Route stdout to a file, if the path is not absolute then the base directory is \$FASTSEARCH/var/log/infotool. Example: infotool -r stdout.log

infotoolConfig.xml

All pipelines are defined in this file.

```
$FASTSEARCH/etc/infotool/InfotoolConfig.xml
```

infotoolConfig.xml example

This file contains configuration of one or more pipelines, where each pipeline contains collectors:

```
<infotool>
  <collectors>
```

```

<default>
    <collector> OperatingSystem      </collector>
    <collector> SystemMapper        </collector>
    <collector> FASTVersion         </collector>
    <collector> LocaleAndTime       </collector>
    <collector> DocumentCounter     </collector>
    <collector> ChecksumGenerator   </collector>
    <collector> CrawlerConfig       </collector>
    <collector> Logs                </collector>
    <collector> ConfigurationReport </collector>
    <collector> InfotoolComplete    </collector>
</default>
.....
</collectors>
</infotool>

```

This configuration specifies a pipeline called `default`. The `default` pipeline name is assigned when no specific pipeline name is provided to infotool via the `-p` option.

The output created by the tool is placed in a directory under `$FASTSEARCH/var/log/infotool`. The name of the directory has the format `infotool_yyyymmdd-HHMMSS`.

Collectors

Infotool invokes various sub-modules, called *collectors*. Some collectors can be controlled by customizing XML files. If a collector is configurable it will have its configuration files in the `$FASTSEARCH/etc/infotool` directory.

InfotoolComplete

Infotool creates a directory `$FASTSEARCH/var/log/infotool/infotool_yyyymmdd-HHMMSS`. This collector makes a zip file, named `infotool-<user@hostname>-<yyymdd>.zip` containing the files from the output directory.

This collector also displays informational text showing what information to provide and how to contact FAST Technical Support.

InfotoolComplete example

All results of running infotool on the host `adminnode.company1.com`, as the user test at August 9th 2007 will be called `infotool-test@adminnode-20070809.zip`.

ChecksumGenerator

This collector traverses the installation and makes md5 sums of the selected files. It is configured by the file `$FASTSEARCH/etc/infotool/ChecksumGenerator.xml`.

ChecksumGenerator example

```

<checksum>
    <files>
        <file>etc/NodeConf.xml</file>
    </files>
    <directories>
        <directory>
            <name>bin</name>
            <exclude><! [CDATA[core.*]]></exclude>
            <exclude><! [CDATA[Workbench\*.jar]]></exclude>
        </directory>
    </directories>

```

```
</directories>
</checksum>
```

This example generates an md5 checksum for the file \$FASTSEARCH/etc/NodeConf.xml and every file under \$FASTSEARCH/bin except those matching core.* and Workbench.jar.

These md5 sums can be used to determine if the versions of files in an installation are consistent, and track changes to the system.

ConfigurationReport

This collector generates structured data based on information gathered by other collectors.

The output of this collector can be found in the file ConfigurationReport.xml and is mainly for use by FAST Technical Support.

CoreDetector

This collector checks the \$FASTSEARCH/bin directory for any core dump files. This collector will not gather any information if used on Windows.

If any core files are detected, the file utility is used to determine which binary generated the core file; date information about the core file is also collected for debugging purposes.

CrawlerConfig

This collector contacts the crawler and retrieves a list of all the collections defined for the crawler. It then retrieves the configuration files for each of the collections.

FASTVersion

This collector tries to determine the version of FAST ESP you are running along with versions of the following components: Crawler, Indexer, QRServer, and Search.

ConfigZipper

This collector creates one zip file, named ConfigZipper.zip by default, that contains the configuration files of the FAST ESP installation. It is controlled by the file \$FASTSEARCH/etc/infotool/ConfigZipper.xml.

ConfigZipper example

```
<filezipper>
  <files>
    <file>InstallProfile.xml</file>
  </files>
  <directories>
    <directory>
      <name>etc</name>
    </directory>
    <directory>
      <name>var/etc</name>
    </directory>
    <directory>
      <name>var/searchctrl</name>
    </directory>
    <directory>
      <name>var/templates</name>
    </directory>
    <directory>
      <name>lib/python2.3/processors</name>
    </directory>
    <directory>
      <name>lib/python2.2/processors</name>
    </directory>
  </directories>
</filezipper>
```

```
</directories>
</filezipper>
```

This configuration collects the contents of the following files:

- \$FASTSEARCH/InstallProfile.xml
- \$FASTSEARCH/etc/
- \$FASTSEARCH/var/etc
- \$FASTSEARCH/var/searchctrl
- \$FASTSEARCH/var/templates
- \$FASTSEARCH/lib/python2.3/processors
- \$FASTSEARCH/lib/python2.2/processors

Any directories that do not exist are silently ignored.

FreeSpace

This collector:

- Traverses all the nodes in the system to calculate the amount of free space on disk.
- Retrieves rtsearchrc.xml from each node and checks for the configured free-space minimum for indexing.
- Compares the free space on each node to verify that it meets the needs of the most demanding cluster configured in rtsearchrc.xml.

Logs

This collector is configured in the file \$FASTSEARCH/infotool/etc/LogsCollector.xml. This file configures a number of collections of log files, which are copied into separate zip files.

Logs example

```
<logscollector>
  <collection name="errors">
    <startpoint>var/log</startpoint>
    <filepattern>*log</filepattern>
    <filepattern type="regexp">.*scrap[\.]*[0]*</filepattern>
    <exclude-path>querylogs</exclude-path>
    <exclude-path>infotool</exclude-path>
    <max-lines>1000</max-lines>
    <max-age>10 days</max-age>
    <match-pattern>ERROR</match-pattern>
    <match-pattern>FAILURE</match-pattern>
  </collection>
</logscollector>
```

This example assumes a configuration file with one collection.

This configuration generates a file called logs-errors.zip. The startpoint parameter determines where the collection will start, if the path is not absolute then the value of the \$FASTSEARCH environment variable is prepended to it. In this case the collection will start at \$FASTSEARCH/var/log.

The files collected must match one of the provided filepattern patterns. In this case they must either end with *log (using regular pattern expansion) or they must match the regular expression .*scrap[\.]*[0]*.

The exclude-path parameters contain subpaths that should not be collected. In this case, \$FASTSEARC/var/log/querylogs is excluded.

There are two ways of limiting the size of the collected file. In this case we specify max-lines which means that only the specified number of lines will be collected from the end of the file. You may also specify max-size, in bytes. That also is counted from the end of the file. If neither is specified, the last 1000 lines are collected.

The parameter `max-age` specifies the maximum age a file may have to be included in the collection. The default value is 10 days. The value is specified in the format `xx(days|weeks|hours)`, for example, 2 weeks.

If the parameter `match-pattern` is present, it should contain strings that are required to be present in the lines that are to be included from the file. In this example, only the lines containing `ERROR` and `FAILURE` are included in the collection.

There is also a parameter named `verbatim-endings` that specifies files that should be included in full, without being checked for patterns or split. The default value of this parameter is `.zip` and `.gz`.

OperatingSystem

This collector tries to determine which operating system and which version you are running.

For Windows, the collector tries to determine system version (2000, 2003, etc.) and service pack level.

For Linux, the collector retrieves `/etc/*release` to determine system distribution and version of the distribution. It also tries to determine the versions of some of the required packages.

SystemMapper

This collector retrieves as much information about the current system layout as possible. It contacts the ConfigServer to obtain the list of nodes, then queries the nodes on what services they are running. This information is useful in itself, but is also used by other collectors (for example, the FreeSpace collector uses it to traverse various nodes for available free space).

HardwareProfile

This collector retrieves information about the hardware of the system, and how it is configured.

It determines:

- Network device configuration
- Software RAID configuration (Linux)
- Physical disk information

DocumentCounter

This collector retrieves a list of collections that are registered on the system from the ConfigServer. It then gathers a document count for each collection from the QRServer.

This collector obtains a document count from each indexer in an installation.

On an idle system, this stage can be used to diagnose problems with orphaned documents.

Eventlog

This collector retrieves entries from the Windows eventlog and stores them in an XML file in the infotool results directory at `$FASTSEARCH/var/log/infotool`.

The collector can be configured to perform filtering on the events by age or by number of events. The configuration of the collector is performed by editing the file `$FASTSEARCH/etc/infotool/Eventlog.xml`.

The following example shows a configuration that collects events less than 10 days old, up to a maximum of 100 events.

Eventlog example

```
<eventlog>
    <max-events>100</max-events>
    <max-age>10 days</max-age>
</eventlog>
```

LocaleAndTime

This collector retrieves a list of collections that are registered on the system from the ConfigServer. It then gathers a document count for each collection from the QRServer.

ViewCollector

This collector retrieves enough information to replicate a Search View in FAST ESP.

The tool `$FASTSEARCH/bin/view-admin` is used to get a list of all Search Views on a system. All published search views are then exported using the `$FASTSEARCH/bin/exportsearchprofile` tool.

The current index-profile is gathered for replication purposes which is necessary for importing a Search View.

Chapter

3

Log file information

Topics:

- [*About log files*](#)
- [*Log file types*](#)
- [*Guidelines for using log files*](#)
- [*Log file information for individual modules*](#)

This chapter provides log file information including where the files are located, what types of files are available, and how to use the files.

About log files

The log files are produced by the Log Server. The log files are by default located in the subdirectories of `$FASTSEARCH/var/log` and can be viewed in the **Logs** selection of the FAST ESP Administrator Interface.

Log file types

There are several types of log files:

- `*.scrap` files contain standard output messages for a specific module
- `*.scrap.0` files contain standard output messages for a specific module from the previous run
- `*.logs` files contain module activity logs
- All log files related to a given module are by convention collected in the subdirectory
`$FASTSEARCH/var/log/<module>`

To adjust log levels for components, change the setting in `$FASTSEARCH/etc/NodeConf.xml`. Most components use the `-l` flag and a hex code to indicate a log level.

Guidelines for using log files

If you have an issue, use the following guidelines to find a possible solution:

- View scrap files in `$FASTSEARCH/var/log`. Each scrap file has a module associated with it. The name of the file corresponds to the module.
- Make sure that processes start when you execute `$FASTSEARCH/nctrl start`.
- View information in the **Logs** selection of the FAST ESP Administrator Interface.

Log file information for individual modules

Log file information at the module level can help isolate an issue.

The following table lists the FAST ESP modules:

Module	Description
WebAnalyzer	Processes links that are retrieved from documents.
Configuration server	Responsible for document processing via an assigned pipeline.
Content API or SDK	
Crawler	Crawls the web to get content.
Document processor	Responsible for document processing via an assigned pipeline.
File traverser	Utility used to feed documents from disk into the system.

Module	Description
License manager	Responsible for validation of FAST ESP license and applying the license to modules that request it. License files that do not have a SERVER line, will allow FAST ESP to be installed without a License Manager. However, if the license does have a SERVER line, FAST ESP will not function without this module. In all cases, FAST ESP will not function without the license file.
Node controller	Controls FAST ESP software on one node.
QRServer	Responsible for receiving queries and performing query transformations if needed.
Real-time search engine/Real-time search indexer	Responsible for search and result processing and indexing.
Storage service	Stores internal data.

Configuration server

Log files:

\$FASTSEARCH/var/log/configserver/configserver.scrap

\$FASTSEARCH/var/log/syslog/configserver_*.log

Usage:

- Check the \$FASTSEARCH/var/log/archive/ folder for the latest zipped file.
- If the error occurred before the current log, then attach the latest zipped all.log.

Content API or SDK

Log files:

\$FASTSEARCH/var/log/syslog/all.log

\$FASTSEARCH/var/log/syslog/contentdistributor_*.log

Usage:

If you need to contact FAST Technical Support regarding a Content API or SDK issue:

1. Send exact snippet messages from the feeder.
2. Send out the code.

Crawler

Refer to the *Enterprise Crawler Guide* for detailed log file information.

Document processor(s)

Log files:

Scrap file for a processor server <x>, where <x> is an integer 1....N:

```
$FASTSEARCH/var/log/procserver/procserver_<x>.scrap
```

Log file for a processor server running on <host> at <port>:

```
$FASTSEARCH/var/log/syslog/procserver_<host>_<port>.log
```

Configuration files:

Configuration file for a custom stages is located in:

```
$FASTSEARCH/etc/processors/
```

Configuration file for pipeline configuration is located in:

```
$FASTSEARCH/etc/processors/PipelineConfig.xml
```

Usage:

To monitor document processing the `psctrl` and `doclog` tools must be used together. The `psctrl` tool is used to turn on additional logging and debugging while the `doclog` tool shows the information about pipelines and stages.

psctrl

Tool: `$FASTSEARCH/bin/psctrl [options] [commands]`

Usage: Refer to the *Operations Guide, psctrl and doclog tools*, for detailed command options.

To run `psctrl`, connect to the configuration server and retrieve statistics about all known document processors.

Use the `-P` option and specify the <host:port> of the document processor you want to access. The port base default is <base_port> + 3200 and will increment by 5 for each document processor process.

For example, with a default port of 13000 you can reach document processors by connecting to the following ports: 16200, 16205, 16210, ...

This mode is the most useful for monitoring because it continues to function even if the configuration server is down.

doclog

Tool: `$FASTSEARCH/bin/doclog [options]`

Usage: Refer to the *Operations Guide, psctrl and doclog tools*, for detailed command options.

psctrl and doclog examples

Example 1 - To turn on document logging and show the most recent log:

```
$FASTSEARCH/bin/psctrl doctrace on  
$FASTSEARCH/bin/doclog -a
```

Example 2 - To turn on document logging and show all document logs with errors:

```
$FASTSEARCH/bin/psctrl doctrace on
$FASTSEARCH/bin/doclog -e
```

Example 3 - To check process status:

```
$FASTSEARCH/bin/psctrl status
```

Example 4 - To turn on document logging, show all document logs with errors and save the output into a file `output.txt`:

```
$FASTSEARCH/bin/psctrl doctrace on
$FASTSEARCH/bin/doclog -e > output.txt
```

Example 5 - To check for missing documents in the index, use the commands in Example 4, then check the Error: information displayed in the `output.txt`.

esp4jtool-dictman dictionary management tool

Log files:

esp4jtool-dictman logs to:

```
$ESP4J_HOME/apps/esp4jtool-dictman[-local]/log/esp4jtool-dictman[-local].log
```

In server mode, the adminserver log provides additional information concerning the adminserver in which the dictionary service runs as well:

```
$FASTSEARCH/adminserver/logs/adminserver.log
```

```
$FASTSEARCH/esp4j/apps/resourceservice/log
```

File traverser

Log files:

Enable fileraverser logging with the `-l` parameter.

```
$FASTSEARCH/var/log/syslog/all.log
```

Usage:

If you need to contact FAST Technical Support regarding a File Traverser issue:

1. Send the exact command or batch files being used for content processing.
2. Cut and paste the output from the fileraverser.
3. Send exact problem description in snapshot or log files.

License manager

Log files:

```
$FASTSEARCH/var/log/lmgrd/lmgrd.scrap
```

```
$FASTSEARCH/var/log/syslog/all.log
```

Lock file (UNIX only):

```
$FASTSEARCH/var/tmp/lockFASTSRCH
```

License file location:

```
$FASTSEARCH/etc/fastsearch.lic
```

Usage:

Use these files to find out if a license is valid or if there are any processing issues.

Node controller**Log files:**

```
$FASTSEARCH/var/log/syslog/nodecontroller_<host>_<port>.log
```

Usage:

Check log file for problems starting or stopping processes and servers using the nctrl start, stop, suspend, resume commands. Refer to the *Operations Guide*, *nctrl tool*, for detailed information.

QRServer**Log files:**

Process related log files:

```
$FASTSEARCH/var/log/qrserver/qrserver.scrap  
$FASTSEARCH/var/log/searchctrl/*.*  
$FASTSEARCH/var/log/syslog/qrserver_*.log  
$FASTSEARCH/var/log/syslog/all.log
```

Query related log files (such as advanced linguistic issue):

```
$FASTSEARCH/var/log/querylogs/*
```

Usage:

To monitor the QRServer, connect it to the server HTTP interface. Perform a simple query using FAST query language and check for proper results. For example:

```
http://<host>:15100
```

on the host that runs the QRServer. Also:

- QRServer configuration is available at <http://:15100/configuration>
- Top-level dispatcher configuration information is available at <http://:15151/status>
- Dispatcher configuration information is available at <http://:15601/status>

Format:

The query log consists of one line per HTTP request on the query interface, where each line consists of several fields which are separated by space. There are two formats, one for search results and one for other HTTP requests on the query interface.

The query line consists of the following fields:

```
<Client-IP> - - <Timestamp> "<HTTP-Request-Line>" <ErrorCode> <Content-Length> "<Referer>"<br/>"<Agent>" <RequestTime> <SearchTime> <DocSumTime> <TotalHits>
```

<SearchTime>, <DocSumTime> and <TotalHits> are only included for queries.

The following table describes the format for the individual fields of each query log line:

Query log field	Description
Client-IP	The IP-address of the HTTP client.
Timestamp	The time of the log entry (not the start of the HTTP request) Example: [18/Jun/2003:11:25:24 +0200]
HTTP-Request-Line	The HTTP request received.
ErrorCode	The HTTP error code. Values below 1000 are standard HTTP error codes (200 is OK). The product specific error codes above 1000 are described in the <i>Query Integration Guide, Using the HTTP query interface</i> .
Content-Length	Number of bytes returned for the HTTP request (the query result returned to the client in XML or text format). When using the Query API Search APIs the actual result is returned according to the text template format as described in the <i>Query Integration Guide, Using the HTTP query interface</i> .
Referer	The value of the Referer HTTP header field.
RequestTime	The total time spent serving the request (<seconds>.<milliseconds>).
SearchTime	The time spent waiting for initial search results from the search engines (<seconds>.<milliseconds>). The split between SearchTime and DocSumTime is related to the 2-phase request process used by the query & result server (QRServer). The SearchTime equals the time to perform the query matching and returning the document IDs and the rank values. The query & result server (QRServer) will then retrieve the result fields (document summaries) from the search engines.
DocSumTime	The time spent waiting for document summaries from the search-subsystem (<seconds>.<milliseconds>). Refer to SearchTime for additional information.
TotalHits	The total number of hits matching the query.

Query log example

```
27.0.0.1 - - [18/Jun/2003:11:25:24 +0200] "GET
/cgi-bin/asearch?hits=10&offset=0&type=all&query=car&filter=%2blanguage:en
HTTP/1.1" 200 506 "" "Mozilla/5.0 (X11; U; Linux i686; en-US; rv:1.4)
Gecko/20030529" 0.3187 0.2542 0.0000 0
```

The HTTP GET query string typically starts with /cgi-bin/<template>? where <template> indicates how the search result is presented in the HTTP response:

search

- asearch: Returns text based results for the default result view.
 - <result-view>: Returns text based results for an alternative result view as defined in the index profile.
- xsearch
- xml-search: Returns XML based results for the default result view.
 - xml-<result-view>: Returns XML based results for an alternative result view as defined in the index profile.

Reporting issues

If you find you need to contact FAST Technical Support regarding a QRServer issue:

1. Send exact product description with snapshot of the **System Management** pages.
2. Include information on any changes including network or FAST ESP configuration changes, that might have contributed to the issue.
3. Send a snapshot of the search view including error messages.

Real-time search engine/Real-time search indexer

Log files:

RTSearch log files on the target node:

```
$FASTSEARCH/var/log/rtsearch/*.log
```

RTSearch configuration files:

```
$FASTSEARCH/etc/config_data/RTSearch/webcluster/rtsearchrc.xml
```

```
$FASTSEARCH/etc/config_data/RTSearch/webcluster/index-profile
```

General system log:

```
$FASTSEARCH/var/log/syslog/all.log
```

Usage:

Storage service

Log files:

```
$FASTSEARCH/var/log/storageservice/storageservice.scrap
```

Usage:

Check for PostgreSQL database usage errors.

WebAnalyzer

Log files:

```
$FASTSEARCH/var/log/webanalyzer/*  
$FASTSEARCH/var/log/fdm/*  
$FASTSEARCH/var/log/fdmworker/*  
$FASTSEARCH/var/log/walinkstorerreceiver/*
```

```
$FASTSEARCH/var/log/walookupdb/*  
$FASTSEARCH/var/log/wapartialupdater/*
```

Usage:

If you need to contact FAST Technical Support regarding a WebAnalyzer issue:

1. Send the all the log file mentioned above.
2. Send \$FASTSEARCH/var/log/syslog/all.log file.

Chapter

4

Known issues and troubleshooting tips

Topics:

- [*Document processor issues*](#)
- [*Lemmatization tips and issues*](#)
- [*Synonyms issues*](#)
- [*Real-time search engine tips and issues*](#)
- [*QRServer issues*](#)
- [*License manager tips and issues*](#)
- [*Disk configuration tips*](#)
- [*Network tips*](#)
- [*SBC issues*](#)

This chapter provides resolutions to known issues and troubleshooting tips.

Document processor issues

Optimize document processing performance

Issue: I want to optimize document processing performance.

Resolution: The `-o` parameter in the document processor section of the `$FASTSEARCH/etc/NodeConf.xml` file allows you to determine whether you want to maximize debugging capabilities or maximize document processing performance.

Pipeline Statistics is by default turned ON (debugging is enabled) to optimize document processing debugging capabilities. Complete the following to turn Pipeline Statistics OFF (debugging is disabled) and optimize document processing performance:

1. Open the document processor section in `$FASTSEARCH/etc/NodeConf.xml`.
2. Edit the file by adding the optimize `-O` switch. The changed configuration should look something like this:

```
<process name="procserver" description="Document Processor" multi="yes">
  <start>
    <executable>procserver.exe</executable>
    <parameters>-b 1 -O -c 0 -P $PORT -ORBInitRef
      NameService=corbaname::myhostname.com:16099</parameters>
```

3. Add a maximum documents log switch `-d 1000` to keep the memory usage at a reasonable level. **This option determines how many documents are kept in the document log.**
4. Save the file.
5. Run the command:
`$FASTSEARCH/bin/nctrl reloadcfg`
6. Restart each Processor Server. This can be done in one of two ways:
 - Select **System Management** from the FAST ESP Administrator Interface and **Restart the Processor Servers**
 - Run the command:
`$FASTSEARCH/bin/nctrl restart procserver_<n>`

Unusual behavior during document processing

Issue: Something unusual happens to my documents when they travel to pipeline.

Resolution:

1. Turn on document debugging as described in **Optimize document processing performance** issue.
2. Inspect the logs. If you notice pipeline errors or are not sure about what each stage does, contact FAST Solution Services.

Documents skipping stages

Issue: Documents skip certain stages in the pipeline.

Resolution: This behavior is normal as some documents may not be suitable for certain stages. This is not an issue.

PDF documents missing attributes in result set

Issue: My PDF documents are missing certain attributes, such as title, in the result set.

Resolution: The PDF extractor reads PDF documents and their attributes as specified during the creation of the PDF files. Go to **File->Document Properties** to verify the properties of all documents in Adobe Acrobat. This list also displays the available document properties and their values. Properties with no values will show up empty (or "null") in the result set.

To perform an additional verification the properties of PDF files:

1. Edit the pipeline used to process the documents by adding an instance of the **Spy** stage to it. Spy describes what document properties were recorded during the process. This is useful if your documents have properties that are not showing up in the result set.
2. Review the results of the **Spy** stage and contact FAST Technical Support if there are any issues.

It is also possible to add a custom document processor that checks whether a certain property, such as title, is empty. If a property is empty, the document processor will create a title based on the body or other attribute. Contact FAST Technical Support to find out how this can be implemented.

Lemmatization tips and issues

Lemmatization tips

If you have problems with lemmatization, check to make sure that:

1. Lemmatization is enabled.
 - Check the index profile to see if the overall installation as well as any fields are enabled.
 - The field you are searching has `lemmatize="yes"`. The composite field you are searching has `lemmas="yes"`. Note that lemmas are only generated for composite fields, not for single fields.
 - Lemmatization is enabled for queries (`flag qtf_lemmatize=yes`).
2. Language is correctly set for the query. Check the query language and compare with languages in `LemmatizationConfig.xml`.
3. Language identifier is included in the document processing pipeline and the language is set manually. If no language is set, then the lemmatizer document processor will not process a document.
4. Dictionaries are successfully loaded. Check the logs at startup of processor server and qrserver to see that they are listed. Logs will display a warning or error message if dictionaries are missing.
5. If tokenization has been altered, make sure that the:
 - Terms entered in the lemmatization dictionaries are in lower-cased lettering. Terms should not contain blanks or punctuation marks.
 - Tokenizer is present in the pipeline and it is listed in the correct order (the tokenizer stage must be listed before the lemmatizer stage).

Add a language for lemmatization

Issue: You want to add a language for lemmatization.

Resolution:

All steps apply to the configuration file `LemmatizationConfig.xml` and should be performed on all nodes.

1. You will need the lemmatization dictionaries for the language you want to install.

Before requesting additional packages, check to see if the required dictionaries are already present in resources/dictionaries/lemmatization. If you need dictionaries for additional languages, contact FAST Solution Services.

2. Add the language as a standard_lemmatizer item, following the example in the LemmatizationConfig.xml configuration file:

```
<standard_lemmatizer language="es" mode="document_expansion" active="yes">
  <lemmas active="yes" parts_of_speech="NA" />
</standard_lemmatizer>
```

Set the mode to query_expansion, document_expansion or reduction, and choose the appropriate automaton.

The parts_of_speech describes the sequence of Nouns (N), Adjectives (A), Verbs (V).

3. Restart qrserver and procserver.
4. Setting up a new language for lemmatization by reduction or by document expansion requires re-processing of all documents. Re-process all documents.

Change lemmatization for a language

Issue: You want to change the lemmatization strategies for a language.

Resolution:

All steps apply to the configuration file LemmatizationConfig.xml and should be performed on all nodes.

1. You will need the lemmatization dictionaries for the language you want to install.

Before requesting additional packages, check to see if the required dictionaries are already present in resources/dictionaries/lemmatization. If you need dictionaries for additional languages, contact FAST Solution Services.

2. Search for the standard_lemmatizer item for the language you want to change by following the example in the configuration file. Set the mode to query_expansion, document_expansion or reduction.
3. Change the parts_of_speech to the desired level.
4. Restart qrserver and procserver.
5. Setting up a new language for lemmatization by reduction or by document expansion requires re-processing of all documents. Re-process all documents.

Set up default query lemmatization

Issue: You want to set up default query lemmatization for multiple languages.

Resolution: This only applies to lemmatization by query expansion or by reduction. For lemmatization by document expansion, there is no query lemmatization.

1. If not already done, define a standard_lemmatizer for each language. (Refer to the previous Lemmatization issues/resolutions.) All languages that are used in such a composite lemmatizer must be configured in the same lemmatization mode.
2. Add the comma separated list of language in the lemmatization tag as default_query_language:

```
<lemmatization default_mode="reduction" default_query_language ="en,de">
```
3. Restart the qrserver.

Synonyms issues

Index-side synonyms do not work

Issue: Synonyms do not work.

Resolution:

- Verify that all the synonym configuration steps have been applied. Check that the `FASTEARCH/resources/dictionaries/synonyms/dp/<lang>_synonyms.aut` file exists. If it is not, recompile using `esp4jtool-dictman`.
- Make sure the **Lemmatization** box is checked in the **Advanced** section in the **Search Front End**.

Index-side synonyms do not work for all documents

Issue: Synonyms do not work for all documents.

Resolution:

- Re-process the old documents in your system whenever you change the synonym.
- Check to make sure configuration changes use correct format and syntax.
- Understand that only documents with the language you have added synonyms for will be affected. Verify that the documents you are missing have the correct language.

Index-side synonyms do not work for all queries

Issue: Synonyms do not work for all queries.

Resolution: Synonyms are not subject to lemmatization. You must add all forms of a synonym to the dictionary. For example, if you have added the synonym `automobiles`, then only that term would result in hits and the singular form of the word `automobile` will not.

Query-side synonyms seen in the system log

Issue: Errors from QT Synonym are seen in the system log.

Resolution: Check the system log after restarting QRServer. The following lists the most common error cases:

- The QRServer terminates immediately upon start up.

This is most likely caused because the QRServer was started while another instance of it was still running. This can be verified and resolved by stopping the QRServer and checking on the process table of your server to see if the QRServer has indeed been stopped. If QRServer is still running, wait until it disappears from the process table before you restart the new instance of QRServer. If the QRServer had to load a lot of large spellchecking or lemmatization dictionaries, it might take several minutes to stop.

- The QRServer starts, but QT Synonym is not part of it.

To verify this:

Enter the QRServer configuration panel usually found on port 15100. Click **Configuration**.

In the **Active Query Transformers** section, check to make sure there is a line for **QT Synonym**.

If there is a line for **QT Synonym**, then QT Synonym was successfully loaded and is part of the QRServer.

If there is no line for **QT Synonym**, then check the log messages in the FAST ESP Administrator Interface to determine what happened while loading QT Synonym. Check the message immediately preceding the `FastQT_Synonym was disabled` message. This message should specify the cause of the failure.

If the message states `FastQT_Synonym couldn't open dictionary`, then check the following file:

```
$FASTSEARCH/etc/config_data/QRServer/webcluster/etc/qrservr/qtf-config.xml
```

Be sure you use `QT_Synonym` with the correct version of the automata (automata compiled with `esp4jtool-dictman`, `dictcompile` or `makeaut` distributed together with FAST ESP).

- The QRServer runs, but instead of producing synonyms or modifying queries, QT Synonym generates error messages in the log such as:

```
Error while looking up "...." in the synonym dictionary...
```

This happens if the synonym automaton does not have the appropriate format. In general, synonym dictionaries have to be built using the `esp4jtool-dictman` command line tool. Building the synonym dictionary by hand with the tool `makeaut` is not recommended.

Real-time search engine tips and issues

Real-time search engine tips

- Refer to the *Installation Guide* for details on system requirements. Use the `/usr/bin/ssh` default for installation. Do not run as root or admin (unless that is your standard level of operation).
- For re-installation or any shutdown/restart of FAST ESP, have a list of all processes (e.g. `Imgrd` = License Manager) that are currently running. Uninstall by removing the entire `FASTSEARCH` directory. Verify that all hanging processes have been stopped and kill any remaining processes with Task Manager (Windows) or `kill -9` (Unix/Linux). Re-install or restart FAST ESP.

Systems unable to index

Issue: Certain Linux systems may be unable to index due to too many open files.

Indexing is failing and the following messages are displayed:

```
Error messages containing the following text segments indicate that the number of  
open files exceeds the configured system limits:
```

- Too many open files, ...
- Could not open xxx, Too many open files

Resolution: To increase the number of open files:

1. Log in as root.

2. Run the command:

```
ulimit -n 8192
```

3. Run the command:

```
/etc/init.d/sshd restart
```

4. Log out as root.

5. Log in as a FAST user on the server.

6. Stop and start the indexer from the command line:

```
$FASTSEARCH/bin/.ncctrl stop indexer  
$FASTSEARCH/bin/.ncctrl start indexer
```

7. Repeat this procedure on all search nodes.

-
8. Contact your system administrator to make this change permanent.

Identifying and recovering from search node failure

Issue: To identify and recover from a search node failure:

Resolution:

- If one or more of the search processes die, this will be reported from the remote controller or the indexer (depending on whether we are using remote or local search nodes). Either way, the search process will be restarted, and this process will be repeated until it is successful.
- If a restart fails repeatedly, it may be that there is a problem with the configuration, input data, or system resources that is blocking the startup of the process.
- Check the relevant rtsearch logs for details. If possible, resolve the problem indicated in the log, and the restart should succeed. If there is a problem with the index, a reindex of the partition displaying the problem should resolve the issue.

QRServer issues

QRServer and fdispatch not connected

Issue: Error code indicates that the channel is not connected. There is no connection between fdispatch and QRServer.

Resolution:

Restart the fdispatch and QRServer processes and then complete the following general debugging procedure:

1. Turn on fnet debugging:

```
http://<qrserverhost>:15100/control?debug.fnet=1
```

2. Issue some queries.

3. Look at the output in qrserver.scrap for output similar to the following:

```
[2004-01-21 09:33:07] INFO : qrserver->fnet:  
events[/s][loop/int/io][967.3/0.0/1.0] packets[/s][r/w][1.0/1.0]  
data[kB/s][r/w][0.03/0.01]
```

If you find a `packets[/s][r/w][0.0/1.0]` message displayed, then the problem is most likely present on your system (zero packets read per second).

4. Debug logging for RTS and searchctrl. To investigate this further, it is useful to have debug logs from two components.

On all search nodes:

```
etc/searchrc-1.xml  
set debuglog="true"
```

On the configuration server node:

```
etc/config_data/RTSearch/webcluster/rtsearchrc.xml  
set debugLog="true"
```

5. Edit `$FASTSEARCH/etc/searchrc-dispatch.xml` on search nodes with topfdiscpatch.

```
set debuglog="true"
```

By default the option is set to false.

6. Edit the files above accordingly. Shut down the system and restart after verifying that all the processes have terminated successfully (`frtsobj.exe`, `fsearchctrl.exe`).

QRServer timed out

Issue: An error message similar to the following is displayed:

```
Error messages containing the following text segments indicate incorrect timeout  
values in query processing:  
- Docsum fetching timed out  
- Timed out while waiting for query result.  
- Timed out while waiting for document summaries
```

Resolution:

This issue may be due to incorrect timeout values in the system To resolve this issue, increase the source.xml, fdispatch.addon and fsearch.addon timeout values in the system. To do this:

1. Update FASTSEARCH/etc/fdispatch.addon with the following values:

```
maxdocsumwait = 80  
maxsearchwait = 70  
maxsocksilent = 120
```

2. Update FASTSEARCH/etc/fsearch.addon with the following value:

```
maxsocksilent = 120
```

3. Update FASTSEARCH/etc/qrserver/sources.xml with the following value:

```
timeout query="60" docsum="70" (the timeout tag)
```

4. Restart rtsearch/qrserver.

License manager tips and issues

License manager tips

- Only one instance of License Manager is allowed to run on a node.
- If License Manager fails to start, the rest of the services will fail as well. Check the log file if this occurs.
- The most common issues are related to incorrect licenses or a lock file (\$FASTSEARCH/var/tmp/lockFASTSRCH).
- If a lock file exists, (a) review your active processes, (b) kill the License Manager (if any) and (c) remove the lock file.
- If you have a problem with a license, contact FAST Technical Support.

Components not initializing

Issue: Components are not coming up correctly or not coming up at all.

Resolution: This may be a license manager issue. Send the following information to FAST:

- Send the etc/fastsearch.lic file.
- Include var/log/lmgrd/lmgrd.scrap and var/log/syslog/all.log.
- Send any network configuration change information.

QRServer processing is limited

Issue: A QRServer processing limitation appears. This may be a license manager issue.

The following error message is displayed:

```
Error: System is overloaded
```

Resolution: Send the following information to FAST:

- A snapshot of `http://<search nodes>:15151/status/.`
- The `$FASTSEARCH/var/log/rtsearch/*.log` file.
- The `$FASTSEARCH/etc/fastsearch.lic` file.
- The `$FASTSEARCH/var/log/lmgrd/lmgrd.scrap` and `$FASTSEARCH/var/log/syslog/all.log`.

Disk configuration tips

Disk errors are reported by the kernel.

- If the drive itself is actually bad, messages will in most cases show up on the system console, `dmesg`, or `/var/log/messages`.
 - Run the `dmesg` command as root. `dmesg` dumps the contents of the kernel log buffer to standard output.
 - Disk errors may also be logged to the file `/var/log/messages`.
- If you have software raids, inspect the status of those raids by looking in the file `/proc/mdstat`.
- If you have hardware raids, run a program that can talk directly with your RAID controller. If you want to force a read on a disk you suspect have problems, use the `dd` command to dump an entire raw device.

For example:

```
dd if=/dev/sdb2 of=/dev/null
```

In another window, use `dmesg` to check if the kernel reports any errors.

Network tips

FAST ESP can be distributed across several nodes so a functional network is an essential part of the system.

Typically, FAST ESP needs the following services configured to function properly:

- DNS - Nodes must have fully qualified domain names assigned to them. If fully qualified domain names are not used, the following symptoms can occur:
 - License manager is unable to start correctly
 - Configuration server produces errors
 - FAST ESP Administrator Interface displays errors

By default, the hostname is assigned to the first network configuration card. If you have multiple NICs you will either have to disable one of them or configure your FAST ESP setup to bind to the interface that has the hostname assigned.

- Nameserver issues configuration
- Open access (no firewalls) - No firewall must exist between the nodes involved in the same FAST ESP cluster.
- Multicasting - Multicasting is used for copying data if an indexer and a corresponding search service reside on different nodes within the same FAST ESP cluster.
- Secure shell support - Used for copying configuration files. Can be used instead of multicasting for copying index data. One user needs to be created on every node in FAST ESP cluster. This user must have a passwordless login (use keys with no phrases).

Query transmission is slow

Issue: Query transmission is slow even though search functionality is fast.

Resolution: Check to make sure the network card is configured to a **full** speed mode and not an auto-detect mode. Auto detection of networking speed may not work correctly.

Machine configuration issues

Refer to the *Installation Guide* for any machine configuration issues including operating system version or patch level, invalid system paths, missing permissions, software conflicts, port assignments.

SBC issues

SBC query reports cannot be generated

Issue: SBC query reports can not be generated due to unavailable logtransformer.

Resolution:

Set up SBC reports to run on demand. To do so, use the queryreportgenerator tool:

1. Create a collection.
2. Create a search profile with the collection. Refer to the *Search Business Center Guide*.
3. Feed documents from the following files to the collection and make queries from the **Search Front End**.

```
esp/components/logtransformer/bin$ ./launch.sh downloader  
esp/components/logtransformer/bin$ ./launch.sh logmerger  
esp/esp4j$ sh bin/queryreportgenerator
```

Upon completion, the SBC will show an updated query report.

Chapter

5

Log levels and log messages

Topics:

- *Logserver configuration*
- *Determining the destination host for log messages*
- *Log levels and log messages for individual modules*

This chapter provides log levels and log message information.

Logserver configuration

By default, FAST ESP saves log files to subdirectories in the `$FASTSEARCH/var/log` directory on the host machine where the logserver process is running. The `NodeConf.xml` file defines this directory path. Periodically, FAST ESP rolls and zips logs into the `$FASTSEARCH/var/log/archive` directory. These default directory values are configured by way of the `logserver -a` and `-d` flags in `NodeConf.xml` file.

 **Caution:** Before you make modifications to the `NodeConf.xml` file, contact FAST Solution Services. This file is automatically generated by the FAST ESP Installation Utility. It is the primary configuration file for the FAST ESP application.

Logserver configuration example

```
<process name="logserver" description="Log Server">
  <start>
    <executable>logserver</executable>
    <parameters>-P $PORT -m 500000 -f $FASTSEARCH/var/log -a
$FASTSEARCH/var/log/archive -z</parameters>
    <port base="3010"/>
    <delay>5</delay>
  </start>
  <outfile>logserver.scrap
  </outfile>
</process>
```

Determining the destination host for log messages

If you have FAST ESP components of distinct FAST ESP instances running on the same host machine, each component records log messages on the host specified by the `$FASTSEARCH/etc/LoggerConfig.xml` file –where `$FASTSEARCH` is each cluster's value of the current FAST ESP installation directory.

Destination host for log messages example

In the following log markup, all components with the environment settings of the current FAST ESP instance send messages to the FAST ESP instance running on `esp1.example.net`.

```
<?xml version="1.0"?>
<LoggerConfig>
  <outputs>
    <output name="server">
      <info name="hostname">esp1.example.net</info>
    </output>
  </outputs>
</LoggerConfig>
```

In FAST ESP, component processes pipe standard output and standard error to a filename defined by the value of the `<outfile>` attribute. If the `<outfile>` value does not include a directory name, the file is created in a subdirectory of `$FASTSEARCH/var/log`, using the file name (minus suffix) as directory name. In the following example, the output will be produced in `$FASTSEARCH/var/log/httpd/httpd.scrap`. If a relative directory name is specified, it is relative to `$FASTSEARCH/var/log`.

```
<process name="httpd" description="Web Server">
  <start>
    <executable>httpd</executable>
    <parameters>-f $FASTSEARCH/etc/httpd/httpd.conf</parameters>
  </start>
  <outfile>httpd.scrap</outfile>
```

```
<writepid>0</writepid>
</process>
```

Log levels and log messages for individual modules

Each log message produced by FAST ESP has an associated log level. Log messages are filtered based on the log level and the module producing the message, and this filtering is configurable. Refer to *Configuring Logging* in the *Operations Guide*.

Admin server log levels and messages

Admin server log levels and log messages.

Log level	Log message	Cause(s)	Action(s)

Configuration server log levels and messages

Configuration server log levels and log messages. By default, the Configuration Server writes log messages to:

```
$FASTSEARCH/var/log/configserver/configserver.scrap
```

Log level	Log message	Cause(s)	Action(s)
CRITICAL	Contact Fast Search & Transfer (http://www.fastsearch.com/) for a new license.	Your license has expired.	Refer to the licensing information listed on the <i>FAST Support</i> page in this guide for FAST contact information.
CRITICAL	Error loading config file: UnicodeError: ASCII encoding error: ordinal not in range (128)	The parser has encountered characters out of range of the expected encoding.	Edit the configuration file and remove those characters.
CRITICAL	Failed to start ConfigServer: error: (226, 'Address already in use')	A program is already using the port the configserver is trying to listening on.	Start the configserver on another port or shut down the program using the one you are trying to use.
CRITICAL	Unable to perform license checkout: Invalid license file syntax	Your license file has an invalid syntax.	Correct the license file syntax.
ERROR	CollectionError: The collection does not exist	The collection name referenced does not exist.	Use the FAST ESP Administrator Interface to check for names of existing collections. If necessary, create the collection.
ERROR	ConfigError: Unable to load configfile: IOError: [Errno 2] No such file or directory: '<filename>'	The configserver process is unable to load one of its configuration files.	Ensure the specified configuration file is situated in the correct directory.

Log level	Log message	Cause(s)	Action(s)
ERROR	DataSourceError: DataSource at ':0' does not exist	The datasource at the specified host:port does not exist.	Register a datasource at the specified host:port.
ERROR	Error saving main configuration file: IOError: [Errno 28] No space left on device	A partition on the file system is full and the configserver cannot save the main configuration file.	Clear some space to allow the configserver to save configuration.  Caution: Stopping the configserver during these conditions may cause information to be lost.
ERROR	Error saving pipeline configuration file: IOError: [Errno 28] No space left on device	The configserver was unable to save modifications made to the pipeline configuration file \$FASTSEARCH/PipelineConfig.xml because there is no space left on the selected partition of the file system.	Clear space to allow the configserver to save configuration.  Caution: Stopping the configserver during these conditions may cause recently added configuration information to be lost.
ERROR	ModuleError: Module <esp-componentname>@<host>:<port> not found	Either the component identified is in a 'down' or 'suspended' state or it is running on a unintended port.	Ensure it is in a 'up' state with \$FASTSEARCH/bin/nctrl sysstatus. Check the port passed in at run time with ps -eaf. Ensure it is restarted on the specified host at the specified port.
ERROR	ModuleError: No Processor Server registered at '<host>: <port>'	No processorserver is running on host:port.	Start the Document Processor (procserver) at host:port.
ERROR	PipelineError: No ProcessorServer supports ALL stages for pipeline <pipeline-name>	You tried creating a pipeline in which no Document Processor supports all stages.	Check the processor stages you included with the pipeline. Ensure that the corresponding <processor-name>.pyc files are located in \$FASTSEARCH/lib/python2.2/processors on the machine where you are added the procserver. If they are not there, copy them from the same directory on a server machine where they are present and restart the procserver.
ERROR	Could not contact the Deployment Manager [...]	The Deployment Manager did not respond in a timely fashion.	Refer to troubleshooting information for the Adminserver.
WARNING	Module <datasearch-component-name> at <host>:<port> not responding, marking as inactive	The module at host:port did not respond to the configserver periodic ping and therefore designated as inactive in NodeState.xml.	Ensure that stale processes are not persisting for the module in question (ps -eaf). If so, kill the processes, and restart the component using nctrl.

Log level	Log message	Cause(s)	Action(s)
WARNING	ping() to <module>@<host>: <port> raised exception - Connect failed: None	Either the component was recently restarted, it is in a 'down' or 'suspended' state, or if running remotely from the configserver there is some latency or connectivity problems on the network.	Look at the log file for error messages related to this component. If warnings persist and the module is running remote to the configserver check your network. Otherwise, check the status of the component with \$FASTSEARCH/bin/nctrl sysstatus.
WARNING	ping() to module at <host>: <port> raised exception - lib.ext.medusa.pyxmlrpc_client.fault: HTTP connection closed unexpectedly	A ping failed because the HTTP connection was shutdown. The specified module might have been restarted. Otherwise, it is in an aberrant state.	Ensure the specified module was not restarted recently. If the module is rejecting or closing out other connections, it might be in an aberrant state or configured improperly. Restart it, and check its \$FASTSEARCH/var/log/<module-name>.scrap file.
WARNING	ping() to module at <host>: <port> raised exception - socket.error: (111, 'Connection refused')	A ping to the module at host:port failed because the connection was refused - possibly because it is gone.	
WARNING	ping() to <module> <host>:@<port> raised exception - socket.error: (10061, 'Connection refused.')	The connection was refused either because the originator does not have the proper permissions to initiate a session or the connection is in an aberrant state.	Check the state of the specified module (\$FASTSEARCH/bin/nctrl sysstatus). Restart the specified module, restart the configserver, and check the modules \$FASTSEARCH/bin/nctrl<module-name>.scrap file if problem persists.
WARNING	ping() to module at <esp-process>@<host>:<port> raised exception - socket. error: (32, 'Broken pipe')	A ping to the module at host:port was not returned by the destination machine.	Check for other warnings related to the specified FAST ESP component.
WARNING	ping() to module at <esp-hostname>:<port> raised exception - Timeout: None	A ping to the module at host:port timed out.	Check the destination host for connectivity (ping) and the network for traffic (netstat).
WARNING	ProcessorServer at _ANY_:_ NUMBER_ has no classes	The processorserver at host:port has no registered classes with the configserver.	Try restarting the processorserver at host:port.
WARNING	Unexpected starttag 'processorid'	An unexpected start-tag was read from the configserver initialization file (CSConfig.xml).	Remove the unexpected start-tag and try again.
			 Caution: CSConfig.xml is automatically generated by the configserver. Do not edit

Log level	Log message	Cause(s)	Action(s)
			this file until you contact FAST Solution Services.

Content distributor log levels and messages

Content distributor log levels and and log messages. By default, the content distributor writes log messages to the FAST ESP Log service section of the FAST ESP Administrator Interface, or in the file named:

\$FASTSEARCH/var/log/syslog/contentdistributor_*.log

Table 1: Content distributor startup / general messages

Log level	Log message	Cause(s)	Action(s)
CRITICAL	Content Distributor (node <id>): Failed to start http server. The port might be in use.	The content distributor is unable to setup its http server using the port specified on the command line. Normally caused by another application using that port. The content distributor will exit after logging this message.	Determine which process is using the content distributor http port (baseport + 3100). Stop that process and restart the content distributor.
ERROR	Failed to create directory for collection info cache	The content distributor is unable to create its collection info cache directory in <esp-dir>/data.	Determine why the content distributor cannot create this directory inside the FAST ESP installation.
ERROR	ContentDistributor could not get memory, exiting.	Memory allocation fails. The content distributor exits after logging this message.	Free up memory or avoid extreme memory usage by the content distributor by using smaller batches or fewer clients.
ERROR	Unable to connect to Configserver	The content distributor is unable to connect to the FAST ESP configuration server. If this is not the first time the content distributor is started, it will use its collection cache until the configuration server is reachable again.	No action necessary. The connection will be restored when the configuration server is reachable again.
ERROR	Failed to resolve <name>. Giving up...	The content distributor fails to resolve/ look up the dispatcher in the next subsystem (typically the indexing dispatcher).	Make sure the dispatcher in the next subsystem is running. Verify that the content distributor does not contain an invalid reference to the next subsystem in its --next-subsystem command line argument.
WARNING	Unable to write node controller pid file "<nct_filename>". This may cause problems with future termination of the service.	These two warnings are logged if the content distributor is unable to write its PID file in the specified directory at startup. This file is used by the node controller to terminate/ stop the content distributor so the failure	Determine why the content distributor is not allowed to write the .pid file inside the FAST ESP installation. Correct the problem and restart the content distributor.

Log level	Log message	Cause(s)	Action(s)
		to write this file may influence the node controllers ability to stop the content distributor.	

Table 2: Content distributor ssl related messages

Log level	Log message	Cause(s)	Action(s)
ERROR	Unable to create SSL context	The content distributor is unable to set up an SSL context when initializing the SSL support.	Report issue to FAST Technical Support.
ERROR	Bad certificate file: <filename>	The certificate in the specified certificate file is bad.	Provide a valid certificate file with matching keys and certificates.
ERROR	Bad private key file: <filename>	The private key in the specified certificate file is bad.	Provide a valid certificate file with matching keys and certificates.
ERROR	Certificate and private key don't match in: <filename>	The private key and the certificate in the specified certificate file do not match.	Provide a valid certificate file with matching keys and certificates.

Table 3: Fast Data Search DS 4.x compatible support messages

Log level	Log message	Cause(s)	Action(s)
ERROR	Failed to look up operation id <operation-id> when converting callback to FDS 4.x callback.	The FAST Data Search 4.x compatibility implementation is unable to find a mapping for a FAST ESP operation id to the corresponding 4.x document id. This causes a document error in the callback to be lost and may cause problems for the client.	Report issue to FAST Technical Support.
ERROR	Unable to find batch for operation id: <operation-id> (callback: " <subsystem-name> . <state>)	The FAST Data Search 4.x compatibility implementation is unable to find a mapping for a set of FAST ESP operation ids to the corresponding 4.x batch id. This causes the callback to be lost and may cause problems for the client.	Report issue to FAST Technical Support.
ERROR	Unable to setup session (The collection was unknown)	Unable to create a feeding session since the specified collection is unknown to the content distributor.	Verify that the collection exists and that it is spelled correctly when specified in the content API.
WARNING	FDS 4.x compat: Unsupported operation ...	Various FAST Data Search 4.x messages logged if the content distributor received a call to a FAST Data Search 4.x method that is not supported.	If this occurs, then the client is using an unsupported FAST Data Search method that needs to be rewritten if it is to be compatible with FAST ESP.

Log level	Log message	Cause(s)	Action(s)
WARNING	contentdistributor::fds4:: callback_servant: failed to report callback for batch <batch-id>	The FAST Data Search 4.x compatibility implementation is unable to report/ send a callback back to the client application. This happens if the client exits before all callbacks are sent.	No action necessary.

Table 4: Master (status queries, remove collection) messages

Log level	Log message	Cause(s)	Action(s)
ERROR	Failed to perform clear_collection for subsystem <subsystem-name> Processing error: id= <error-code>, message= <description>	These two warnings are logged if the content distributor receives an error in a callback for a clear collection operation.	Report issue to FAST Technical Support.
ERROR	clear collection operation failed for collection <collection-name>, format_error:<description>	Content distributor receives a format_error exception when trying to send a clear collection operation.	Report issue to FAST Technical Support.
WARNING	get_status() failed since collection <collection-name> is unknown	Document status query fails because the specified collection is unknown in the system.	Verify that the collection exists and that it is spelled correctly when specified in the content API.
WARNING	get_status() failed: <description>	Document status query fails with a middleware-related exception. May be an indication that the next subsystem in the feeding pipeline is down. If the indexing dispatcher or the indexer is down you might get this warning.	Verify that the components used when feeding content are up and running.
WARNING	get_status() failed: unknown error	Document status query fails with any other error/exception.	Report issue to FAST Technical Support.
WARNING	Master: unable to register dispatcher. The dispatcher seems to be down/unavailable	Communication with a content distributor dispatcher node fails during initialization. This is only relevant for multi-node content distributor setups.	No action necessary.
WARNING	clear collection operation failed 20 times for collection <collection-name>, giving up...	A clear collection operation has failed 20 times when trying to clear/remove a collection. Content distributor will give up the operation, but the FAST ESP system retries the remove collection operation in a while.	No action necessary. The Deployment Manager will retry this operation.

Log level	Log message	Cause(s)	Action(s)
WARNING	clear collection: timed out waiting for callbacks from subsystems when clearing collection <collection-name>	Content distributor has not received the required callbacks for a clear collection operation within one hour from the time the operation was sent. This may happen for very large collections.	No action necessary. The Deployment Manager will retry this operation.

Table 5: Session related messages

Log level	Log message	Cause(s)	Action(s)
ERROR	session <id>: Failed to create session with component registered as <name>	Content distributor fails to create a feeding session with the next subsystem (typically the indexing dispatcher).	Make sure the dispatcher in the next subsystem is running. Verify that the content distributor does not contain an invalid reference to the next subsystem in its --next-subsystem command line argument.
ERROR	session <id>: op_set[<first-id>,<last-id>]: Failed to update status: Format Error: <description>	Content distributor fails to update the document status for failed feeding operations.	Report issue to FAST Technical Support.
ERROR	session <id>: Failed to send op_set[<first-id>,<last-id>] after <# retries> attempts. giving up	Content distributor has multiple failures when dispatching a batch to processor servers. The batch will be discarded and a callback with errors is sent to the client.	No action necessary. If this behavior persists, report issue to FAST Technical Support.
WARNING	session <id>: op_set[<first-id>,<last-id>]: Failed to update status: No resources: <description>. Will continue to try...	A status update is rejected by the next subsystem with a no_resources exception.	No action necessary. The operation is retried until it succeeds.
WARNING	session <id>: Failed to close session in next subsystem: <subsystem-name>	Content distributor fails to close the session in the next subsystem. This happens when a session is being closed, either because the client has called close() or the session has timed out. The cause is most probably that the dispatcher in the next subsystem is not running.	No action necessary.
WARNING	session <id>: op_set[<first-id>,<last-id>] Failed to send operation set to processor server <procserver-id> Reason: <description>	Receiving a system_resource_error exception from the procserver when sending a batch to it. The reason is generally that the processor server has run out of resources while receiving the batch.	No action necessary. If the problem persists, determine why the processor server fails to receive the operation set and correct the problem with the processor server.

Log level	Log message	Cause(s)	Action(s)
WARNING	<pre>session <id>: Failed to send op_set[<first-id>,<last-id>] to processor server <procserver-id>: system_exception: <description> (Processor server may be down or have been restarted " "recently?) session <id>: Failed to send op_set[<first-id>,<last-id>] to processor server <procserver-id>: middleware_exception: <description> (Processor server may be down or have been restarted " "recently?)</pre>	<p>These two messages are logged if you get either a system_exception or a middleware_exception when sending a document batch to a processor server. This is typically caused by a failure of the processor server leaving an invalid reference to it in the content distributor.</p>	<p>No action necessary.</p> <p>The operation will be retried. If it fails multiple times the following error message will be logged and a callback is given to the client:</p> <pre>session <id>: Failed to send op_set [<first-id>,<last-id>] after <# retries> attempts. giving up</pre>

Table 6: Procserv related messages

Log level	Log message	Cause(s)	Action(s)
WARNING	<pre>Processor server <procserver -id> not responding while processing batch. <first-op- id> - <last-op-id></pre>	<p>Processor server stops responding while processing a batch. The batch will be lost and a callback with errors is issued back to the client by the content distributor.</p>	<p>No action necessary.</p> <p>If the problem persists, determine what causes the processor server to fail while processing the set of documents.</p>

Crawler error codes and messages

Refer to the *Enterprise Crawler Guide* for crawler error codes and messages.

File traverser log levels and messages

File traverser log levels and log messages. The file traverser has no default log file. It is user defined.

Log level	Log message	Cause(s)	Action(s)
ERROR	<pre>Unable to write to .ft.collection.last: reason</pre>	<p>The file traverser was unable to write to the specified file.</p> <p>The file contains the time and date of when it was last run for this particular collection.</p>	<p>Examine the specified reason and take appropriate actions.</p> <p>The most frequent cause of this error is file/directory permission errors (verify that the user has write access) and other file system error conditions such as disk full.</p>
ERROR	<pre>Unable to create File Traverser data directory path</pre>	<p>The file traverser was not able to create the data directory specified.</p>	<p>Ensure the user running the file traverser has write access to the specified directory.</p>

Log level	Log message	Cause(s)	Action(s)
ERROR	Contents of .ft.collection.last is not a valid timestamp value	The file may have been corrupted or truncated on disk.	Delete the file and restart the file traverser. This causes the file traverser to be run on all files, rather than incrementally, as it does not know the last time it was run.
ERROR	Error parsing date specification	The specified date specification does not match any of the supported formats.	Refer to the <i>File Traverser Guide</i> for supported date specification formats.
ERROR	Illegal date specification	The specified date specification does not match any of the supported formats.	Refer to the <i>File Traverser Guide</i> for supported date specification formats.
ERROR	Use plural form in date specification	The specified date specification does not match any of the supported formats.	Example: Use 5 minutesago; do not use 5 minute ago.
ERROR	token in datespec is not recognized as a token	The specified date specification does not match any of the supported formats.	The date specification was detected but an unknown token was specified. Supported tokens are second(s), minute(s), hour(s), day(s) and years(s).
ERROR	Error while processing files: reason	The file traverser failed to successfully process some or all of the files.	This is a generic error message. Examine the specified <i>reason</i> and take appropriate actions.
ERROR	ERROR while removing files: reason	The file traverser failed to successfully process the deletes of some or all files.	This is a generic ERROR message. Examine the specified <i>reason</i> and take appropriate actions.
ERROR	Error while flushing queues: reason	An error occurred while trying to send the last remaining data batches to FAST ESP. This was most likely due to a content feeding problem.	This is a generic error message. Examine the specified <i>reason</i> and take appropriate actions. If the problem persists it may be necessary to examine logs from other components to identify the root cause.
ERROR	Error while shutting down: reason	An error occurred during shutdown of the file traverser. This was most likely due to a content feeding problem.	This is a generic error message. Examine the specified <i>reason</i> and take appropriate actions. If the problem persists it may be necessary to examine logs from other components to identify the root cause.
WARNING	Unable to decode start path to unicode	The option to force the use of UTF-8 as the file system encoding was set and the file	Disable the force UTF-8 option.

Log level	Log message	Cause(s)	Action(s)
		traverser failed to convert from the native file system encoding.	
WARNING	Bad format in filename:lineno	The mapping filename contains a syntax error on the specified line.	Correct the syntax and re-run the file traverser.
WARNING	File/URL name redefined in filename:lineno	The filename or URL was defined more than once in the mapping file.	Ensure all entries in the mapping file are listed only once.
WARNING	Unable to get deleted page		
WARNING	No URL specified for filename in mappingfile	The mapping file did not contain the filename to URL mapping for the specified filename.	Ensure the mapping file contains a map for all files that will be processed.
WARNING	Unable to stat filename	The file traverser was unable to retrieve the file status from the file system for the specified file.	Ensure that the user executing the file traverser has the appropriate file system access, and verify that there are no errors on the file system. Manually open the file to verify.
WARNING	Unable to set doclocation. Dropping file filename		
WARNING	Could not determine when this application was run last. Running it on all files.	The file traverser uses a file on disk to keep track on when the file traverser was last run. This enables it to only process files modified since the previous run. This warning indicates it was unable to read/locate this information.	No action necessary. The file traverser will perform a full traversal. Ensure that the file containing the last run time is not accidentally deleted.
WARNING	The file traverser has waited the maximum configured time (number seconds) for outstanding callbacks, and will exit (this timeout can be configured by using the -Y switch).	Document processing of the last document batches has taken longer than the configured callback timeout. The file traverser normally waits for these callbacks to be able to give the user feedback on documents that failed in processing. When this occurs some callbacks will not be reported from the file traverser.	The callback timeout can be increased for subsequent runs to have the file traverser wait longer. However, if quick file traverser shutdown is needed this may not be desirable and the warning can be ignored.
WARNING	Unexpected result when flushing all batches, now exiting		
WARNING	It is recommended that you use either http:// , https:// or file:// as		

Log level	Log message	Cause(s)	Action(s)
	the scheme for the -p parameter		
WARNING	It is strongly recommended that you do not use UNC paths as the -p parameter		
WARNING	It is strongly recommended that you do not use local drive letters in the -p parameter		

License error messages

Error messages produced by a licensing issue.

Severity	Log message	Cause(s)	Action(s)
[no severity level]	License failed. Shutdown forced.	The license has expired or the license server has been unreachable for several days. Note: The indexer will log a failure and shut down if, after 4 days, it has been able to start but has not been able to properly check out a license. The error message is displayed and nctrl attempts to restart. Upon a restart and still no license being available, the indexer will keep restarting.	Check the admin node and the license for the indexer. If the server is up and running, then a new license should be installed for reactivation. Refer to the licensing information listed on the <i>FAST Support</i> page in this guide for FAST contact information.

Postprocess error codes and messages

Refer to the *Enterprise Crawler Guide* for postprocess error codes and messages.

Node controller log levels and messages

Node controller log levels and and log messages.

By default, the node controller writes messages to:

\$FASTSEARCH/var/log/nctrl/nctrl.scrap

Log level	Log message	Cause(s)	Action(s)
CRITICAL	Process <fds-process-name> has reached its minimum disk space. Shutting down.	Low on disk space.	Free up some disk space on the host machine where the specified process runs.

Log level	Log message	Cause(s)	Action(s)
ERROR	Error starting process: rtsearch: OSerror: [Errno 12] Not enough space	Failed to start Real-Time Search due to an out of memory situation.	Add more RAM and/or increase swap size.
ERROR	Process <fds-process-name> still running after waiting 30 seconds	The specified process did not shutdown when requested.	Ensure it has stopped with either ps -eaf or \$FASTSEARCH/bin/nctrl sysstatus. If it continues running shutdown the process with kill.
ERROR	Process <fds-process-name> was not running, restarting it	The process <process-name> was not running (and it should be) so the Node Controller is attempting to restart it.	Examine logs, look for core files, etc to determine why the process was down.
WARNING	<fds-process-name> terminated by signal SIGBUS	The specified process crashed with UNIX signal SIGBUS.	The Node Controller will attempt to restart the process. If this repeats, contact FAST Technical Support.
WARNING	<fds-process-name> terminated by signal SIGSEGV	The specified process crashed with UNIX signal SIGSEGV (segmentation fault).	The Node Controller will attempt to restart the process. If this repeats, contact FAST Technical Support.
WARNING	ConfigServerError: Failed to register with ConfigServer: Fault: (146, 'Connection refused')	The configserver is not running or the Node Controller is not able to access it.	Verify that the configserver is running, that the contents of etc/CSLocation.xml are correct and that there is network connectivity between the nodes.
WARNING	lmgrd returned with status 243	Return code reported when process exited.	Check license manager (lmgrd) return code documentation for cause and remedy.

Chapter

6

Processes overview

Topics:

- *Processes descriptions*

This chapter lists all processes running in an installation.

Processes descriptions

The following table provides a brief description of all processes running in an installation. Some processes are only running on a subset of the hosts in a multi-host installation, and some processes may not be running at all depending on the system configuration.

Process	System component	Description
bliss	Administration server	This process runs each time the index profile is updated. It manages the deployment of the new configuration specified by the index profile.
cachemanager	Cache manager	Cache manager process.
caster	Indexer	Indexer process used for sending binary indexes to search servers not running on the same host as the indexer. This process will always be running even when there are no remote search servers.
catcher	Search	Search server process that listens for updated binary indexes sent from remote indexer servers. This process is only running if the search server depends on a remote index server.
clarity	Clarity Agent	The agent process that backs the Clarity web GUI.
configserver	Configuration server	Configuration server process.
contentdistributor	Content distributor	Content distributor process.
ESPexec.exe	Node Controller	Windows only wrapper processes used by the Node Controller for process control.
FASTSRCH	License manager	FAST specific process that handles license validation. This is the vendor specific daemon process.
FASTESP.exe	Windows service daemon	Windows only process that invokes the service FAST ESP.
fdispatch	Search	Part of the search server. Accepts incoming queries and dispatches these to the different fdispatch processes. It also accepts the results and merges them. This process runs in two capacities, low level and top level. Low level dispatch runs on each machine with fsearch processes and communicates with the fsearch processes themselves. Top level dispatch runs only when you have multiple columns. Top level dispatch communicates with each Low level dispatch and aggregates results before passing them to the qrserver.
fdmworker	WebAnalyzer	Does the link processing in the WebAnalyzer service. Can also act as an internal file server for the WebAnalyzer.
findex	Indexer	This process handles the last part of an indexing cycle. It manipulates binary index structures to accommodate new or updated content and removes deleted content.
fixmlindex	Indexer	This process handles the first part of an indexing cycle. It parses new documents and translates them into internal binary index structures.

Process	System component	Description
frtsobj	Indexer	This process manages the Indexing server covering indexing itself such as providing APIs for operations on content.
fsearch	Search	This is the core search process. It accepts queries from fdispatch and returns matches from an index corresponding to a search server partition and column.
fsearchctrl	Search	This is the search server control process. It manages the other search processes.
httpd (Apache.exe)	Administration server	The administration user interface is implemented as a set of server-side scripts on this HTTP server.
java	Admin Server	A web server running administrative services and various web GUIs, as FAST Home, the Search Business Center, and the Search Front End.
java	Resource Service	The launcher process for the Resource Service Component.
java	Resource Service	The Resource Service process.
java	Logtransformer	The launcher process for the Logtransformer process.
java	Logtransformer	The Logtransformer scheduler.
lmgrd	License manager	License manager process. Handles license evaluation and routing to vendor specific daemon process.
logserver	Log server	Log server process.
nctrl	Node controller	This process manages starting, stopping and monitoring of all processes on a single host in the system.
omniNames		Internal name service keeping track of host and port bindings for the processes making up the system.
postgresql	Storage server	Storage server process.
procserver	Document processor server	Document processor process.
qrserver	QRServer	This part of the QRServer performs query and result transformations. It uses the QRServer fdispatch process to execute the actual queries.
snmpd	SNMP Agent	SNMP Agent for FAST ESP.
walinkstorerreceiver	WebAnalyzer	Server that receives link data from the procservers and writes it to disk for later processing.
walookupdb	WebAnalyzer	Server that serves out processed link data to the procservers.
webanalyzer	WebAnalyzer	Configuration and scheduling process of the WebAnalyzer service.

